

Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) EP 1 055 989 A1

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
29.11.2000 Bulletin 2000/48

(51) Int Cl.7: G06F 1/00

(21) Application number: 99304164.9

(22) Date of filing: 28.05.1999

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

• Balacheff, Boris
Keynsham, Bristol BS31 2JH (GB)
• Chen, Ligu
Bradley Stoke, Bristol BS32 9DQ (GB)

(71) Applicant: Hewlett-Packard Company
Palo Alto, California 94304-1112 (US)

(74) Representative: Lawman, Matthew John Mitchell
Hewlett-Packard Limited,
IP Section,
Building 3,
Filton Road
Stoke Gifford, Bristol BS34 8QZ (GB)

(72) Inventors:
• Poudler, Graeme John
Stoke Gifford, Bristol BS34 8XQ (GB)

(54) System for digitally signing a document

(57) The preferred embodiment of the invention comprises a computer system which employs a trusted display processor (260), which has a trusted processor (300) and trusted memory (305, 315, 335, 345) physically and functionally distinct from the processor and memory of the computer system. The trusted display processor (260) is immune to unauthorised modification or inspection of internal data. It is physical to prevent forgery, tamper-resistant to prevent counterfeiting, and

has crypto functions (340) to securely communicate at a distance. The trusted display processor (260) interacts with a user's smartcard (122) in order to extract and display a trusted image, or seal (1000), generate a digital signature of the bitmap of a document image and control the video memory (315) so that other processes of the computer system cannot subvert the image during the signing process. The user interacts with the trusted display processor via a trusted switch (135).

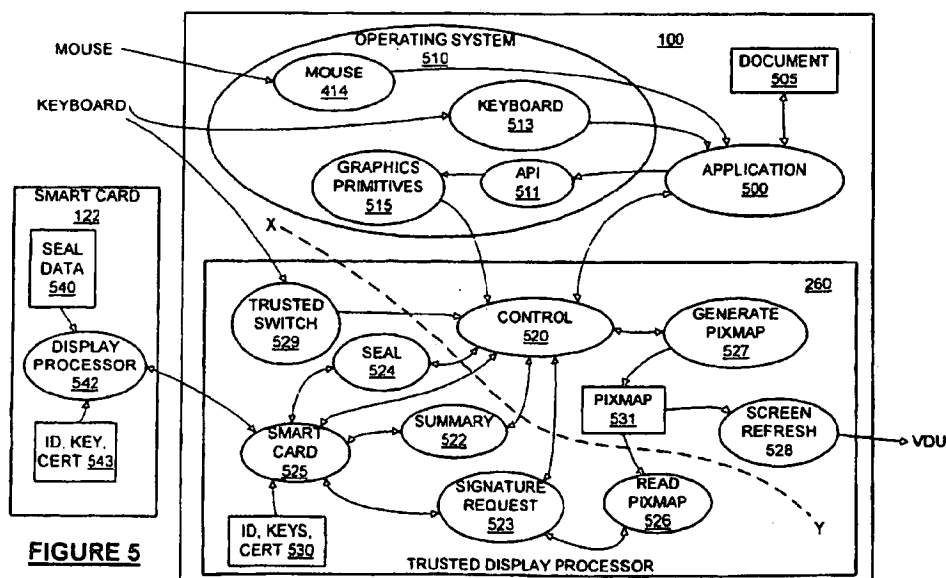


FIGURE 5

EP 1 055 989 A1

Description**Technical Field**

5 [0001] The present invention relates to apparatus and methods for digitally signing image data, and in particular documents, in a manner which provides a high level of confidence to a signing party that the document they think they are signing is in fact the document they are signing.

Background Art

10 [0002] Conventional prior art mass market computing platforms include the well-known personal computer (PC) and competing products such as the Apple Macintosh™, and a proliferation of known palm-top and laptop personal computers. Generally, markets for such machines fall into two categories, these being domestic or consumer, and corporate. A general requirement for a computing platform for domestic or consumer use is a relatively high processing power, Internet access features, and multi-media features for handling computer games. For this type of computing platform, 15 the Microsoft Windows™ 95 and 98 operating system products and Intel processors, so-called WinTel platforms, dominate the market.

[0003] On the other hand, for business use, there are a plethora of available proprietary computer platform solutions available aimed at organizations ranging from small businesses to multi-national organizations. In many of these applications, a server platform provides centralized data storage, and application functionality for a plurality of client stations. For business use, other key criteria are reliability, remote access, networking features, and security features. For such platforms, the Microsoft Windows NT 4.0™ operating system is common, as well as the UNIX and, more recently, the Linux operating systems.

20 [0004] Windows-type operating systems allow a user to run separate applications in separate windows, and provide a so-called WIMP (windows, icons, menus and pointers) interface, whereby a user typically interacts with applications using a keyboard to enter data and a mouse to select options and control applications via dialog boxes and drop-down (or pull-up) menus.

[0005] With the increase in commercial activity transacted over the Internet, known as "e-commerce", there has been much interest in the prior art on enabling data transactions between computing platforms, over the Internet. In particular, 25 it is perceived to be important for users to be able to enter into binding contracts over the Internet, without the need for the current standard hand-signed paper contract. However, because of the potential for fraud and manipulation of electronic data, in such proposals, fully automated transactions with distant unknown parties on a wide-spread scale as required for a fully transparent and efficient market place have so far been held back. The fundamental issue is one of trust between users and their computer platforms, and between interacting computer platforms, for the making of such transactions.

30 [0006] There have been several prior art schemes which are aimed at increasing the security and trustworthiness of computer platforms. Predominantly, these rely upon adding in security features at the application level, that is to say the security features are not inherently embedded in the kernel of operating systems, and are not built in to the fundamental hardware components of the computing platform. Portable computer devices have already appeared on the market which include a smartcard, which contains data specific to a user, which is input into a smartcard reader on the computer. Presently, such smartcards are at the level of being add-on extras to conventional personal computers, and in some cases are integrated into a casing of a known computer. Although these prior art schemes go some way to improving the security of computer platforms, the levels of security and trustworthiness gained by prior art schemes may be considered insufficient to enable widespread application of automated transactions between computer plat- 35 forms. Before businesses expose significant value transactions to electronic commerce on a widespread scale, they will require greater confidence in the trustworthiness of the underlying technology.

[0007] In the applicant's co-pending patent applications 'Trusted Computing Platform' 99301100.6, filed at the European Patent Office on 15 February 1999 and 'Computing Apparatus and Methods of Operating Computing Apparatus' 9905056.9, filed at the UK Patent Office on 5 March 1999, the entire contents of which are incorporated herein by reference, there is disclosed a concept of a 'trusted computing platform' comprising a computing platform which has a 'trusted component' in the form of a built-in hardware component. Two computing entities each provisioned with such a trusted component may interact with each other with a high degree of 'trust'. That is to say, where the first and second computing entities interact with each other the security of the interaction is enhanced compared to the case where no trusted component is present, because:

- 55 • A user of a computing entity has higher confidence in the integrity and security of his own computer entity and in the integrity and security of the computer entity belonging to the other party;

- Each entity is confident that the other entity is in fact the entity which it purports to be;
- Where one or both of the entities represent a party to a transaction, e.g. a data transfer transaction, because of the in-built trusted component, third party entities interacting with the entity have a high degree of confidence that the entity does in fact represent such a party;
- The trusted component increases the inherent security of the entity itself, through verification and monitoring processes implemented by the trusted component; and
- The computer entity is more likely to behave in the way it is expected to behave.

[0008] As has been indicated above, the conventional method of signing a document is to physically write a signature on the medium (usually paper) upon which an image of a document is reproduced. This method has the advantages that it is clear what is being signed, and the signed image is proof of what was signed. However, it does not meet the needs of e-commerce.

[0009] Nowadays it is also possible to digitally sign a document, using a conventional computer platform and standard encryption techniques. In conventional computer platforms, however, the present inventors have appreciated that the electronic rendition of a document which is digitally signed is typically not the same rendition of the document that is visible to the user. It is therefore possible for a user to unintentionally sign data that is different from that which he intended to sign. Conversely, it is also possible for a user to intentionally sign data and later fraudulently claim that the signed data does not correspond to that displayed to him by the computer platform. Such problems would still be the present, even if a trusted platform, as described above, were used.

[0010] Conventional electronic methods of signing are well known to those skilled in the art. Essentially, digital data is compressed into a digest, for example by the use of a hash function. Then that digest is encrypted by the use of some encryption method that has been initialised by a secret key (or simply a 'secret'). This is normally done on a computer platform, such as a PC. One implementation is to sign data using a private encryption key held secret on a user's smartcard, which is plugged into a smartcard reader attached to the computer platform. In the specific case of a textual document, the digital data may be the file produced by a word processor application, such as Microsoft's Notepad, Wordpad, or Word. As usual, the act of signing implies that the signer accepts some legal responsibility for the meaning of the data that was signed.

[0011] Hash functions are well-known in the prior art and comprise one way functions which are capable of generating a relatively small output data from a relatively large quantity of input data, where a small change in the input data results in a significant change in the output data. Thus, a data file to which is applied a hash function results in a first digest data (the output of the hash function). A small change e.g. a single bit of data in the original data file will result in a significantly different output when the hash function is reapplied to the modified data file. Thus, a data file comprising megabytes of data may be input into the hash function and result in a digital output of the order of 128 to 160 bits length, as the resultant digest data. Having a relatively small amount of digest data generated from a data file stored in the reserved directory is an advantage, since it takes up less memory space and less processing power in the trusted component.

[0012] During known signing processes, a user will typically interpret a document as it has been rendered on the computer's monitor at normal magnification and resolution. In existing applications, the user's smartcard signs data in a format that is the representation of the document by the application used to create and/or manipulate the document. The present inventors believe, however, that there is potential for software to send data to the smartcard that has a different meaning from that understood by the user when viewing the screen. This possibility may be sufficient reason to introduce doubt into the validity of conventional methods of digitally signing electronic representations of documents that are to be interpreted by people.

Disclosure of the Invention

[0013] The invention consists of systems and methods to improve confidence in digitally signed documents that are to be interpreted by people. They necessarily involve the reliable display of data, which can be used for other purposes.

[0014] In accordance with a first aspect, the present invention provides a data processing system arranged to generate a digital signature representative of a document, the data processing system comprising:

- main memory means for storing a document to be digitally signed;
- main processing means for executing at least one application process comprising means to generate graphics signals for displaying the document;
- means for generating a request signal for the document to be signed;

a display system comprising:

frame buffer memory;

means for generating digital image data representative of the document on the basis of the graphics signals and storing the digital image data in the frame buffer memory; and

5 means for reading the digital image data from the frame buffer memory, converting the data into signals suitable for displaying an actual image thereof on a display means and forwarding said signals to a display means; and

a trusted component comprising independent processing means operable, in response to receipt of the request signal, for generating a digital signature representative of the digital image data.

10 [0015] As such, the digital signature is generated on the basis of the data used to display the document to a user.

[0016] In preferred embodiments, the trusted component comprises means for denying to any unauthorised application or process write access to at least the portion of the frame buffer memory containing the digital image data of the document, and means for generating a digital signature representative of the digital image data while the respective portion of the frame buffer memory is not accessible for writing data to.

15 [0017] In preferred embodiments the data processing system further comprises a removable token comprising processing means for receiving the digital image data, or a representation thereof, from the frame buffer memory and generating a respective digital signature. Conveniently, the removable token is an appropriately programmed smart-card.

20 [0018] In preferred embodiments the trusted component comprises means for acquiring and/or generating trusted image data and means for controlling the display system to highlight the displayed document image using the trusted image data. This provides visual feedback to a user that the trusted component is in control of the operation.

[0019] The trusted image data may comprise pixmap data representative of the trusted image or instructions for forming the trusted image.

25 [0020] Preferably, the trusted component comprises means for acquiring and/or generating trusted image data from a removable token.

[0021] The trusted component preferably further comprises means for controlling the display system to display messages to a user.

30 [0022] In preferred embodiments, the data processing system further comprises trusted input means by which a user can respond to messages in a secure fashion. The trusted input means may comprise a switch connected to the trusted component via a secure communications channel.

[0023] In preferred embodiments, the trusted component and the secure token enact a mutual authentication process in advance of further interactions.

35 [0024] In preferred embodiments, the trusted component forms an integral part of the display system. For example, the display system may be arranged such that the trusted component is physically and functionally positioned between the main processing means and the frame buffer memory, such that the main processing means can only access the frame buffer memory indirectly through functions of the trusted component.

[0025] Preferably, the trusted component further comprises means for generating data summarising a digital signature operation.

40 [0026] Other aspects and embodiments of the invention will become apparent from the following description, claims and drawings.

Brief Description of the Drawings

45 [0027] Embodiments of the present invention will now be described in detail with reference to the accompanying drawings, of which:

Figure 1 is a diagram which illustrates a computer system suitable for operating in accordance with the preferred embodiment of the present invention;

50 Figure 2 is a diagram which illustrates a hardware architecture of a host computer suitable for operating in accordance with the preferred embodiment of the present invention;

Figure 3 is a diagram which illustrates a hardware architecture of a trusted display processor suitable for operating in accordance with the preferred embodiment of the present invention;

Figure 4 is a diagram which illustrates a hardware architecture of a smart card processing engine suitable for operating in accordance with the preferred embodiment of the present invention;

55 Figure 5 is a diagram which illustrates a functional architecture of a host computer including a trusted display processor and a smart card suitable for operating in accordance with the preferred embodiment of the present invention;

Figure 6 is a flow diagram which illustrates the steps involved in generating an individual signature of a document;

Figure 7 is diagram which illustrates the sequence of messages between the trusted display processor and the smart card in order to recover seal image data from the smart card;

Figure 8 is diagram which illustrates the sequence of messages between the trusted display processor and the smart card in order to generate a signature of a document image;

Figure 9 is diagram which illustrates the sequence of messages between the trusted display processor and the smart card in order to generate a signature of a summary of the document image signing process;

Figure 10a is a diagram which illustrates an exemplary trusted image;

Figures 10b to 10d are diagrams which illustrate the visual steps in signing a document image; and

Figures 10e to 10g are diagrams which illustrate alternative ways of highlighting the image of a document to be signed.

Best Mode For Carrying Out the Invention, & Industrial Applicability

[0028] The preferred embodiment utilises a trusted component that most conveniently uses some of the characteristics of the 'trusted component' described in the applicant's co-pending European patent application number 99301100.6. In that application, the trusted component is a hardware device, comprising a processor programmed to measure an integrity metric of its host computer, compare it with a true value of the integrity metric and communicate the integrity (or otherwise) of the host computer to users or other host computers. The significant similarities between that trusted component and the trusted component in the preferred embodiment herein are:

that they both use cryptographic processes but preferably do not provide an external interface to those cryptographic processes;

that they are both tamper-resistant or tamper-detecting, so that their operation cannot be subverted, at least without the knowledge of the legitimate user; and

that they both preferably consist of one physical hardware component that is both physically and functionally independent of the host computer on which it resides.

[0029] Such independence is achieved by the trusted component having its own processing capability and memory.

[0030] Techniques relevant to tamper-resistance are well known to those skilled in the art of security, as described in the applicant's co-pending application. These techniques include methods for fabricating components to resist tampering, methods for detecting tampering, and methods for eliminating data when tampering is detected. It will be appreciated that, although tamper-proofing is a most desirable feature of the present invention, it does not enter into the normal operation of the invention and, as such, is beyond the scope of the present description.

[0031] In this description, the term 'trusted', when used in relation to a physical or logical component or an operation or process, implies that the behaviour thereof is predictable under substantially any operating condition and highly resistant to interference or subversion by external agents, such as subversive application software, viruses or physical interference.

[0032] The term 'host computer' as used herein refers to a data processing apparatus having at least one data processor, at least one form of data storage and some form of communications capability for interacting with external entities, such as peripheral devices, users and/or other computers locally or via the Internet. The term 'host computer system' in addition to the host computer itself includes standard external devices, such as a keyboard, mouse and VDU, that attach to the host computer.

[0033] The term 'document', as used herein, includes any set of data that can be visualised using a host computer system. Commonly a document will be a textual document, such as a contract. However, a document may comprise graphics, or pictures, instead of, or as well as, text. In general, a document may comprise a single page or multiple pages.

[0034] The term 'pixmap', as used herein, is used broadly to encompass data defining either monochrome or colour (or greyscale) images. Whereas the term 'bitmap' may be associated with a monochrome image only, for example where a single bit is set to one or zero depending on whether a pixel is 'on' or 'off', 'pixmap' is a more general term, which encompasses both monochrome and colour images, where colour images may require up to 24 bits or more to define the hue, saturation and intensity of a single pixel.

[0035] As will become apparent, the trusted component according to the preferred embodiment herein provides a secure user interface and, in particular, controls at least some of the display functionality of its host computer. The trusted component herein may or may not also acquire integrity metrics according to the trusted component in applicant's co-pending patent application, although such acquisition of integrity metrics will not be considered herein.

[0036] In essence, the preferred embodiment enables a user to digitally sign a document stored on a host computer using the private key of the user's smartcard, or other form of secure token such as a cryptographic co-processor. The signing is enacted by a trusted display processor (i.e. the trusted component) of the host computer under conditions

that provide the user with a high level of confidence that the document being viewed on screen is in fact the document the smartcard is signing. In particular, the smartcard carries trusted image data, or a 'seal', which is passed to the host computer over a secure channel and displayed by the trusted component during the signing procedure. It is in part the display of the trusted image, which is typically unique to the user, which provides the user with the confidence that the trusted component is in control of the signing operation. In addition, in the preferred embodiment, the host computer provides a trusted input device, connected directly to the trusted display processor, by which the user can interact with the host computer in a manner which cannot be subverted by other functions of the host computer.

[0037] More particularly, the trusted display processor or a device with similar properties is associated with video data at a stage in the video processing beyond the point where data can be manipulated by standard host computer software. This allows the trusted display processor to display data on a display surface without interference or subversion by the host computer software. Thus, the trusted display processor can be certain what image is currently being displayed to the user. This is used to unambiguously identify the image (pixmap) that a user is signing. A side-effect of this is that the trusted display processor may reliably display any of its data on the display surface, including, for example, the integrity metrics of the prior patent application, or user status messages or prompts.

[0038] Figure 1 illustrates a host computer system according to the preferred embodiment, in which the host computer is a Personal Computer, or PC, which operates under the Windows NT™ operating system. According to Figure 1, the host computer 100 is connected to a visual display unit (VDU) 105, a keyboard 110, a mouse 115 and a smartcard reader 120, and a local area network (LAN) 125, which in turn is connected to the Internet 130. Herein, the smartcard reader is an independent unit, although it may be an integral part of the keyboard. In addition, the host computer has a trusted input device, in this case a trusted switch 135, which is integrated into the keyboard. The VDU, keyboard, mouse, and trusted switch can be thought of as the human/computer interface (HCI) of the host computer. More specifically, the trusted switch and the display, when operating under trusted control, as will be described, can be thought of as a 'trusted user interface'. Figure 1 also illustrates a smartcard 122 for use in the present embodiment as will be described.

[0039] Figure 2 shows a hardware architecture of the host computer of Figure 1.

[0040] According to Figure 2, the host computer 100 comprises a central processing unit (CPU) 200, or main processor, connected to main memory, which comprises RAM 205 and ROM 210, all of which are mounted on a motherboard 215 of the host computer 100. The CPU in this case is a Pentium™ processor. The CPU is connected via a PCI (Peripheral Component Interconnect) bridge 220 to a PCI bus 225, to which are attached the other main components of the host computer 100. The bus 225 comprises appropriate control, address and data portions, which will not be described in detail herein. For a detailed description of Pentium processors and PCI architectures, which is beyond the scope of the present description, the reader is referred to the book, "The Indispensable PC Hardware Handbook", 3rd Edition, by Hans-Peter Messmer, published by Addison-Wesley, ISBN 0-201-40399-4. Of course, the present embodiment is in no way limited to implementation using Pentium processors, Windows™ operating systems or PCI buses.

[0041] The other main components of the host computer 100 attached to the PCI bus 225 include: a SCSI (small computer system interface) adaptor connected via a SCSI bus 235 to a hard disk drive 240 and a CD-ROM drive 245; a LAN (local area network) adaptor 250 for connecting the host computer 100 to a LAN 125, via which the host computer 100 can communicate with other host computers (not shown), such as file servers, print servers or email servers, and the Internet 130; an IO (input/output) device 225, for attaching the keyboard 110, mouse 115 and smartcard reader 120; and a trusted display processor 260. The trusted display processor handles all standard display functions plus a number of further tasks, which will be described in detail below. 'Standard display functions' are those functions that one would normally expect to find in any standard host computer 100, for example a PC operating under the Windows NT™ operating system, for displaying an image associated with the operating system or application software. It should be noted that the keyboard 110 has a connection to the IO device 255, as well as a direct connection to the trusted display processor 260.

[0042] All the main components, in particular the trusted display processor 260, are preferably also integrated onto the motherboard 215 of the host computer 100, although, sometimes, LAN adapters 250 and SCSI adapters 230 can be of the plugin type.

[0043] Figure 3 shows a preferred physical architecture for the trusted display processor 260. In accordance with the preferred embodiment, the trusted display processor 260 is a single hardware component having the characteristics of a trusted component, providing the standard display functions of a display processor and the extra, non-standard functions for generating digital signatures and providing a trusted user interface. The skilled person will appreciate that the functions could alternatively be physically split into two or more separate physical components. However, it will be appreciated on reading the following description that integration of all functions into a single trusted component provides a most elegant and convenient solution.

[0044] According to Figure 3, the trusted display processor 260 comprises:

a microcontroller 300;

non-volatile memory 305, for example flash memory, containing respective control program instructions (i.e. firmware) for controlling the operation of the microcontroller 300 (alternatively, the trusted display processor 260 could be embodied in an ASIC, which would typically provide greater performance and cost efficiency in mass production, but would generally be more expensive to develop and less flexible);

an interface 310 for connecting the trusted display processor 260 to the PCI bus for receiving image data (i.e. graphics primitives) from the CPU 200 and also trusted image data from the smartcard 122, as will be described; frame buffer memory 315, which comprises sufficient VRAM (video RAM) in which to store at least one full image frame (a typical frame buffer memory 315 is 1-2 Mbytes in size, for screen resolutions of 1280x768 supporting up to 16.7 million colours);

a video DAC (digital to analogue converter) 320 for converting pixmap data into analogue signals for driving the (analogue) VDU 105, which connects to the video DAC 320 via a video interface 325;

an interface 330 for receiving signals directly from the trusted switch 135;

volatile memory 335, for example DRAM (dynamic RAM) or more expensive SRAM (static RAM), for storing state information, particularly received cryptographic keys, and for providing a work area for the microcontroller 300;

a cryptographic processor 340, comprising hardware cryptographic accelerators and/or software, arranged to provide the trusted display processor 260 with a cryptographic identity and to provide authenticity, integrity and confidentiality, guard against replay attacks, make digital signatures, and use digital certificates, as will be described in more detail below; and

non-volatile memory 345, for example flash memory, for storing an identifier I_{DP} of the trusted display processor 260 (for example a simple text string name), a private key S_{DP} of the trusted display processor 260, a certificate $Cert_{DP}$ signed and provided by a trusted third party certification agency, such as VeriSign Inc., which binds the trusted display processor 260 with a signature public-private key pair and a confidentiality public-private key pair and includes the corresponding public keys of the trusted display processor 260.

[0045] A certificate typically contains such information, but not the public key of the CA. That public key is typically made available using a 'Public Key Infrastructure' (PKI). Operation of a PKI is well known to those skilled in the art of security.

[0046] The certificate $Cert_{DP}$ is used to supply the public key of the trusted display processor 260 to third parties in such a way that third parties are confident of the source of the public key and that the public key is a part of a valid public-private key pair. As such, it is unnecessary for a third party to have prior knowledge of, or to need to acquire, the public key of the trusted display processor 260.

[0047] The trusted display processor 260 lends its identity and trusted processes to the host computer and the trusted display processor has those properties by virtue of its tamper-resistance, resistance to forgery, and resistance to counterfeiting. Only selected entities with appropriate authentication mechanisms are able to influence the processes running inside the trusted display processor 260. Neither an ordinary user of the host computer, nor any ordinary user or any ordinary entity connected via a network to the host computer may access or interfere with the processes running inside the trusted display processor 260. The trusted display processor 260 has the property of being "inviolable".

[0048] Originally, the trusted display processor 260 is initialised with its identity, private key and certificate by secure communication with the trusted display processor 260 after it is installed onto the motherboard of the host computer 100. The method of writing the certificate to the trusted display processor 260 is analogous to the method used to initialise smartcards by writing private keys thereto. The secure communications is supported by a 'master key', known only to the trusted third party (and to the manufacturer of the host computer 100), that is written to the trusted display processor 260 during manufacture, and used to enable the writing of data to the trusted display processor 260. Thus, writing of data to the trusted display processor 260 without knowledge of the master key is not possible.

[0049] It will be apparent from Figure 3 that the frame buffer memory 315 is only accessible by the trusted display processor 260 itself, and not by the CPU 200. This is an important feature of the preferred embodiment, since it is imperative that the CPU 200, or, more importantly, subversive application programs or viruses, cannot modify the pixmap during a trusted operation. Of course, it would be feasible to provide the same level of security even if the CPU 200 could directly access the frame buffer memory 315, as long as the trusted display processor 260 were arranged to have ultimate control over when the CPU 200 could access the frame buffer memory 315. Obviously, this latter scheme would be more difficult to implement.

[0050] A typical process by which graphics primitives are generated by a host computer 100 will now be described by way of background. Initially, an application program, which wishes to display a particular image, makes an appropriate call, via a graphical API (application programming interface), to the operating system. An API typically provides a standard interface for an application program to access specific underlying display functions, such as provided by Windows NT™, for the purposes of displaying an image. The API call causes the operating system to make respective graphics driver library routine calls, which result in the generation of graphics primitives specific to a display processor, which in this case is the trusted display processor 260. These graphics primitives are finally passed by the CPU 200

to the trusted display processor 260. Example graphics primitives might be 'draw a line from point x to point y with thickness z' or 'fill an area bounded by points w, x, y and z with a colour a'.

[0051] The control program of the microcontroller 300 controls the microcontroller to provide the standard display functions to process the received graphics primitives, specifically:

5

receiving from the CPU 200 and processing graphics primitives to form pixmap data which is directly representative of an image to be displayed on the VDU 105 screen, where the pixmap data generally includes intensity values for each of the red, green and blue dots of each addressable pixel on the VDU 105 screen;

10

storing the pixmap data into the frame buffer memory 315; and
periodically, for example sixty times a second, reading the pixmap data from the frame buffer memory 315, converting the data into analogue signals using the video DAC and transmitting the analogue signals to the VDU 105 to display the required image on the screen.

15

[0052] Apart from the standard display functions, the control program includes a function to mix display image data received from the CPU 200 with trusted image data to form a single pixmap. The control program also manages interaction with the cryptographic processor and the trusted switch 135.

20

[0053] The trusted display processor 260 forms a part of the overall 'display system' of the host computer 100; the other parts typically being display functions of the operating system, which can be 'called' by application programs and which access the standard display functions of the graphics processor, and the VDU 105. In other words, the 'display system' of a host computer 100 comprises every piece of hardware or functionality which is concerned with displaying an image.

25

[0054] As already mentioned, the present embodiment relies on interaction between the trusted display processor 260 and the user's smartcard 122. The processing engine of a smartcard suitable for use in accordance with the preferred embodiment is illustrated in Figure 4. The processing engine comprises a processor 400 for enacting standard encryption and decryption functions, to support digital signing of data and verification of signatures received from elsewhere. In the present embodiment, the processor 400 is an 8-bit microcontroller, which has a built-in operating system and is arranged to communicate with the outside world via asynchronous protocols specified through ISO 7816-3, 4, T=0, T=1 and T=14 standards. The smartcard also comprises non-volatile memory 420, for example flash memory, containing an identifier I_{SC} of the smartcard 122, a private key S_{SC} , used for digitally signing data, and a certificate $Cert_{SC}$, provided by a trusted third party certification agency, which binds the smartcard with public-private key pairs and includes the corresponding public keys of the smartcard 122 (the same in nature to the certificate $Cert_{DP}$ of the trusted display processor 260). Further, the smartcard contains 'seal' data SEAL in the non-volatile memory 420, which can be represented graphically by the trusted display processor 260 to indicate to the user that a process is operating securely with the user's smartcard, as will be described in detail below. In the present embodiment, the seal data SEAL is in the form of an image pixmap, which was originally selected by the user as a unique identifier, for example an image of the user himself, and loaded into the smartcard 122 using well-known techniques. The processor 400 also has access to volatile memory 430, for example RAM, for storing state information (such as received keys) and providing a working area for the processor 400, and an interface 440, for example electrical contacts, for communicating with a smart card reader.

30

[0055] Seal images can consume relatively large amounts of memory if stored as pixmaps. This may be a distinct disadvantage in circumstances where the image needs to be stored on a smartcard 122, where memory capacity is relatively limited. The memory requirement may be reduced by a number of different techniques. For example, the seal image could comprise: a compressed image, which can be decompressed by the trusted display processor 260; a thumb-nail image that forms the primitive element of a repeating mosaic generated by the trusted display processor 260; a naturally compressed image, such as a set of alphanumeric characters, which can be displayed by the trusted display processor 260 as a single large image, or used as a thumb-nail image as above. In any of these alternatives, the seal data itself may be in encrypted form and require the trusted display processor 260 to decrypt the data before it can be displayed. Alternatively, the seal data may be an encrypted index, which identifies one of a number of possible images stored by the host computer 100 or a network server. In this case, the index would be fetched by the trusted display processor 260 across a secure channel and decrypted in order to retrieve and display the correct image. Further, the seal data could comprise instructions (for example PostScript™ instructions) that could be interpreted by an appropriately programmed trusted display processor 260 to generate an image.

35

[0056] Figure 5 shows the logical relationship between the functions of the host computer 100, the trusted display processor 260 and the smartcard 122, in the context of enacting a trusted signing operation. Apart from logical separation into host computer 100, trusted display processor 260 or smartcard 122 functions, the functions are represented independently of the physical architecture, in order to provide a clear representation of the processes which take part in a trusted signing operation. In addition, the 'standard display functions' are partitioned from the trusted functions by a line x-y, where functions to the left of the line are specifically trusted functions. In the diagram, functions are repre-

40

45

50

55

sented in ovals, and the 'permanent' data (including the document image for the duration of the signing process), on which the functions act, are shown in boxes. Dynamic data, such as state data or received cryptographic keys are not illustrated, purely for reasons of clarity. Arrows between ovals and between ovals and boxes represent respective logical communications paths.

5 [0057] In accordance with Figure 5, the host computer 100 includes: an application process 500, for example a wordprocessor process, which requests the signing of a document; document data 505; an operating system process 510; an API 511 process for receiving display calls from the application process 500; a keyboard process 513 for providing input from the keyboard 110 to the application process 500; a mouse process 514 for providing input from the mouse 115 to the application process 500; and a graphics primitives process 515 for generating graphics primitives on the basis of calls received from the application process via the API 511 process. The API process 511, the keyboard process 513, the mouse process 514 and the graphics primitives process 515 are build on top of the operating system process 510 and communicate with the application process via the operating system process 510.

10 [0058] The remaining functions of the host computer 100 are those provided by the trusted display processor 260. These functions are: a control process 520 for co-ordinating all the operations of the trusted display processor 260, and for receiving graphics primitives from the graphics primitives process and signature requests from the application process 500; a summary process 522 for generating a signed summary representative of a document signing procedure in response to a request from the control process 520; a signature request process 523 for acquiring a digital signature of the pixmap from the smartcard 122; a seal process 524 for retrieving seal data 540 from the smartcard 122; a smartcard process 525 for interacting with the smartcard 122 in order to enact challenge/response and data signing tasks required by the summary process 522, the signature request process 523 and the seal process 524; a read pixmap process 526 for reading stored pixmap data 531 and passing it to the signature request process 523 when requested to do so by the signature request process 523; a generate pixmap process 527 for generating the pixmap data 531 on the basis of graphics primitives and seal image data received from the control process 520; a screen refresh process 528 for reading the pixmap data, converting it into analogue signals and transmitting the signals to the VDU 105; and a trusted switch process 529 for monitoring whether the trusted switch 135 has been activated by the user. The smartcard process 525 has access to the trusted display processor's identity data I_{DP} , private key S_{DP} data and certificate $Cert_{DP}$ data 530. In practice, the smart card and the trusted display processor interact with one another via standard operating system calls.

15 [0059] The smartcard 122 has: seal data 540; a display processor process 542 for interacting with the trusted display processor 260 to enact challenge/response and data signing tasks; smartcard identity data I_{SC} , smartcard private key data S_{SC} and smartcard certificate data $Cert_{SC}$ 543.

[0060] A preferred process for signing a document using the arrangement shown in Figures 1 to 5 will now be described with reference to the flow diagram in Figure 6.

20 [0061] Initially, in step 600, the user controls the application process 500 to initiate a 'signature request' for digitally signing a document. The application process 500 may be realised as a dedicated software program or may be an addition, for example a macro, to a standard word processing package such as Microsoft's Word. In either case, neither the signature request nor the application process 500 need to be secure. When the user initiates the signature request, he also specifies the document to be signed, if it is not one which is already filling the whole screen. For example, the document may be displayed across a part of the full screen area or in a particular window. Selection of a particular area on screen is a simple task, which may be achieved in several ways (using a WIMP environment), for example by drawing a user-defined box bounding the area or by simply specifying co-ordinates.

25 [0062] Next, in step 602, the application process 500 calls the control process 520 to sign the image that is being displayed (within a defined area or window) on the screen; the control process 520 receives the call. In parallel, although it is not shown, the control process 520 receives any graphics primitives from the graphics primitives process and forwards them onto the generate pixmap process 527. The call from the application process 500 to sign a document includes the co-ordinates (a,b,c,d) of the edges of the document. Note that this sending of co-ordinates generally enables the signing of the entire surface of the screen, a complete window, or of an arbitrary part of the screen. The application process 500 then waits for the control process 520 to return the signature of the image.

30 [0063] In response to the signature request, in step 604, the control process 520 forces the image that is to be signed to be 'static' from the time of the request until the process has been completed. Herein, 'static' means that the document image cannot be modified other than by the trusted display processor 260. This is so that the user can be certain that what he sees is what he is signing at all times during the process. In the present embodiment, the control process 520 achieves a 'static' display by 'holding-off, or not processing, any further graphics primitives. In some situations, the graphics primitives process (or equivalent) may 'buffer' graphics primitives until the control process 520 is ready to receive further graphics primitives. In other situations, graphics primitives for the image to be signed may simply be lost. Where the document image fills the whole screen, making the image static is simply a case not processing any graphics primitives. However, where the image to be signed forms only a subset, for example a window, of the full screen, the control process 520 needs to determine whether received graphics primitives would affect the 'static' area,

and reject ones that would. As such, the pixmap of the static document image in the frame buffer memory 315 remains unchanged by any instructions from the graphics primitives process, or any other process executing on the CPU 200, while the document image is static.

[0064] Once the document image has been made static, in step 606, the control process 520 instructs the generate pixmap process 527, including in the call the co-ordinates (a,b,c,d) provided by the application process 500, to modify the pixmap to highlight the document to be signed, as will be described in more detail below with reference to Figure 10c. Then, in step 608, if a smartcard 122 is not already inserted in the smartcard 122 reader 120, as determined by the smart card process 525, the control process 520 instructs the generate pixmap process to display a graphical message asking the user to insert his smartcard 122. This message is accompanied by a ten second countdown timer COUNT. If the countdown timer expires (i.e. reaches zero), as a result of not receiving the smartcard 122, the control process cancels the signing operation in step 614 and returns an exception signal to the application process 500. In response, the application process 500 displays an appropriate user message in step 616. If the smartcard 122 is inserted in time, or is already present, then the process continues.

[0065] Next, in step 618, the control process 520 calls the seal process 524, and the seal process 524 calls the smartcard process 525, to recover the seal data 540 from the smartcard 122. Optionally, the control process 520 calls the generate pixmap process 527 to display another message indicating to the user that recovery of the seal data 540 is being attempted. In steps 618 and 620, the smartcard process 525 of the trusted display processor 260 and the display processor process 542 of the smartcard 122 interact using well known, 'challenge/response' techniques to enact mutual authentication and pass the seal data 540 from the smartcard and back to the control process 520. The details of the mutual authentication process and passing of the seal data 540 will now be described with reference to Figure 7.

[0066] According to Figure 7, the smartcard process 525 sends a request REQ1 to the smartcard 122 to return the seal data SEAL 540. The display processor process 542 generates a nonce R_1 and sends it in a challenge to the smartcard process 525. The smartcard process 525 generates a nonce R_2 and concatenates it with nonce R_1 , signs the concatenation $R_1 || R_2$ with its private key to produce a signature $s_{DP}(R_1 || R_2)$, and returns the concatenation $R_1 || R_2$, the signature $s_{DP}(R_1 || R_2)$ and the certificate $Cert_{DP}$ back to the display processor process 542 of the smartcard 122. The display processor process 542 extracts the public key of the trusted display processor 260 from the certificate $Cert_{DP}$, and uses this to authenticate the nonce R_1 and the signature $s_{DP}(R_1 || R_2)$ by comparison with the concatenation $R_1 || R_2$, to prove that the seal request came from the expected trusted display processor 260 and that the trusted display processor 260 is online.

[0067] The nonces are used to protect the user from deception caused by replay of old but genuine signatures (called a 'replay attack') by untrustworthy processes.

[0068] The display processor process 542 of the smartcard 122 then concatenates R_2 with its seal data SEAL 540, signs the concatenation $R_2 || SEAL$ using its private key S_{SC} to produce a signature $s_{SC}(R_2 || SEAL)$, encrypts the seal data SEAL 540 using its private key S_{SC} to produce encrypted seal data 540 $s_{SC}(SEAL)$, and sends nonce R_2 , the encrypted seal data $s_{SC}(SEAL)$, the signature $s_{SC}(R_2 || SEAL)$ and the smartcard's certificate $Cert_{SC}$ to the smartcard process 525 of the trusted display processor 260. The smartcard process 525 extracts the smartcard's public key from the certificate $Cert_{SC}$ and uses this to verify nonce R_2 and the signature $s_{SC}(R_2 || SEAL)$, decrypt the seal data SEAL 540 from the encrypted seal data 540 $s_{SC}(SEAL)$ and, finally, return the seal data SEAL 540, via the seal process 524, to the control process 520.

[0069] Returning to Figure 6, in step 622, when the control process 520 receives the seal data SEAL 540, it forwards the data to the generate pixmap process 527, and instructs the generate pixmap process 527 to generate a seal image and use it to highlight the document to be signed, as will be described below with reference to Figure 10d. Then, in step 624, the control process 520 instructs the generate pixmap process 527 to display a message to the user asking whether they wish to continue with the signing operation. This message is accompanied by a ten second countdown timer COUNT. If the countdown timer expires, in step 626, as a result of not receiving a response from the user, the control process cancels the signing operation, in step 628, and returns an exception signal to the application process 500. In response, the application process 500 displays an appropriate user message in step 629. If, in step 630, the user responds positively by actuating the trusted switch 135 within the ten second time limit, the process continues. The authorisation to continue could alternatively be supplied over an unreliable channel, rather than by using a trusted switch 135, or even using appropriate software routines, providing a reasonable level of authentication is used. Alternatively, it may be decided that the mere presence of an authentic smartcard may be sufficient authorisation for the signing to occur. Such alternatives are a matter of security policy.

[0070] Next, in step 632, the control process 520 instructs the signature request process 523 to request the signing of the document image; the signature request process 523 calls the read pixmap process 526 to request return of a digest of the pixmap data of the document to be signed; and the read pixmap process 526 reads the respective pixmap data, uses a hash algorithm to generate a digest D_{PIX} of the pixmap data and returns the digest to the signature request process 523. Additionally, the read pixmap process 526 generates 'display format data' FD, which includes information

necessary to reconstruct the image from the pixmap data into a text-based document at a later time (FD is not essential, since the document text may not need to be reconstructed), and returns this also to the signature request process 523. For example, the display format data FD may include the number of pixels on the screen surface and their distribution, such as '1024 by 768', and the font type and size used for the text (if the document is text-based) in the document (at least some of this information may instead, or in addition, be contained in a document 'summary', as will be described below). In steps 634 and 636, the signature request process 523 interacts with the display processor process 542 of the smartcard 122 using well-known challenge/response processes to generate an individual signature of the document, as will now be described in detail with reference to the flow diagram in Figure 8.

[0071] According to Figure 8, the smartcard process 525 generates a request REQ2 for the smartcard 122 to generate a signature of the digest D_{PIX} and display format data FD. The display processor process 542 of the smartcard 122 responds by generating a nonce R_3 and sending it to the smartcard process 525 with a challenge to return the digest D_{PIX} and the display format data FD. The smartcard process 525 concatenates the digest D_{PIX} with the display format data FD and nonce R_3 , and signs the concatenation $D_{PIX} || FD || R_3$ to produce a signature $s_{DP}(D_{PIX} || FD || R_3)$. The smartcard process 525 then sends the concatenation $D_{PIX} || FD || R_3$ and its respective signature $s_{DP}(D_{PIX} || FD || R_3)$ to the display processor process 542 of the smartcard 122. The display processor process 542 uses the trusted display processor's public key (which it has already received in the seal data 540 exchange) to verify the trusted display processor's signature $s_{DP}(D_{PIX} || FD || R_3)$ and nonce R_3 , to prove that the digest is the current image digest. The display processor process 542 signs the digest of the pixmap D_{PIX} and the display format data FD, using its private key, to produce two signatures $s_{SC}(D_{PIX})$ and $s_{SC}(FD)$ respectively. The display processor process 542 of the smartcard then returns the signed digest $s_{SC}(D_{PIX})$ and signed display format data $s_{SC}(FD)$ to the smartcard process 525 of the trusted display processor 260. The smartcard process 525 next verifies the digest D_{PIX} and display format data FD, using the smartcard's public key (which it already has as a result of the seal data 540 exchange), and verifies the smartcard's signature, to prove that the smartcard is still online.

[0072] Returning to Figure 6, in step 638, the smartcard process 525 of the trusted display processor 260 concatenates the pixmap PIX, the smartcard's signed versions of the pixmap digest $s_{SC}(D_{PIX})$ and display format data $s_{SC}(FD)$ to form an individual signature $PIX || s_{SC}(D_{PIX}) || s_{SC}(FD)$ of the image, and returns it, via the signature request process 523, to the control process 520, which returns the individual signature to the application process 500. The application process 500 stores the individual signature, in step 640, and responds with a further call to the control process 520 to 'summarise the signing' operation in step 642. The purpose of a summary is to complete the signature, as will be described with reference to the flow diagram in Figure 9 and also the example summary below:

```

1  TC-88503-00.01
2  Access time: Thu 06-May-1999,11:18
35 3  Pages: 2

4  Image01 | 560 x 414 (187,190) [1024 x 768]
5  -----BEGIN SIGNATURE-----
6  PmftitUGoW2h6SLDgqQAvGZ2Y47Fp8wx5ZqE5HS8bGrSV3RD7LKw0kyXPY6yhGDpVNUc/R2
+Gr4mm0LqS/twYuPdskyL4uk3no0w3LG2+f+/vzC4cKMPeY/LhbazzScvhK3CJ+apQxyikj
40 cY5rTC563klovOPTBI/Iyq2PxRnic=
7  -----END SIGNATURE-----

8  Image02 | 670 x 379 (201,228) [1024 x 768]
9  -----BEGIN SIGNATURE-----
15 10 UV1w5Rgr5F0iAjuUW4GP28NKAA+tOy42uBbP78JeQ5w20MI1afTYkSNtfn9VykyMPIfZLwM
7ZZV+4fFttuSgOZ14n5iBkSEwtEj0z6ik/np6paq+01GQZhhJCbq8OaX97Gmdg3AoBq4x+d
hujmqkCJO+Dz6+x8kE24Z8YFXLPOI=
11 -----END SIGNATURE-----

12 Summary signature:
13 -----BEGIN SIGNATURE-----
50 14 ciZDZL2+41FsFci2cPjWfSfltkyXrfHBUM1kAEyudaZcVxD3XczTN7txSazInM2deJL9qnA
een2DW1ZGjplEESNkhoZXj0kT5TYNv2y1YfK01SN+JVF09bmc9GdYLo/hSOWyYG/U29Mzqz
ktaTdTqY/gPh1GajrSJGqRms+we/c=
15 -----END SIGNATURE-----

```

[0073] In step 644, the control process 520 calls the summary process 522 to generate a summary message SUM containing the number of images (two in the example summary above) plus the individual signatures of the images (lines 6 and 10 of the example), a label identifying the trusted display device (line 1 in the example), the current time

and date (line 2 in the example) and a signed digest of the summary itself (line 14 in the example). For each image, the summary also includes the size of the image in pixels (e.g. 560x414 for image 1), the offset from the origin of the screen in pixels (e.g. 187, 190 for image 1) and the display resolution in pixels (e.g. 1024x768 for image 1).

[0074] The summary process 522 then generates a digest of the summary D_{SUM} in step 646 and calls the smartcard process 525 of the trusted display processor 260 to interact with the smartcard 122 using challenge/response processes to generate a signature of the summary digest D_{SUM} , as will now be described with reference to Figure 9.

[0075] According to Figure 9, the smartcard process 525 generates a request REQ3 for the smartcard 122 to generate a signature of the digest of the summary D_{SUM} . The display processor process 542 of the smartcard generates a nonce R_4 and sends it in a challenge to return the digest of the summary D_{SUM} . The smartcard process 525 concatenates the digest D_{SUM} with nonce R_4 and signs the concatenation $D_{SUM}||R_4$ to produce a signature $s_{SDP}(D_{SUM}||R_4)$. The smartcard process 525 of the trusted display processor 260 then sends the concatenation $D_{SUM}||R_4$ and respective signature $s_{SDP}(D_{SUM}||R_4)$ to the display processor process 542. The display processor process 542 then verifies the trusted display processor's signature and nonce R_4 , using the trusted display processor's public key (which it already has from the seal data 540 exchange), to prove that the summary is the current summary. Next, the display processor process 542 signs the digest of the summary D_{SUM} using its private key and sends the signed digest $s_{SC}(D_{SUM})$ to the smartcard process 525. The smartcard process 525 of the trusted display processor 260 verifies the digest and verifies the smartcard's signature, using the smartcard's public key, to prove that the smartcard is still online.

[0076] Returning to Figure 6, in step 652, the smartcard process 525 returns the summary SUM concatenated with the signed digest of the summary $s_{SC}(D_{SUM})$ (to form concatenation $SUM||s_{SC}(D_{SUM})$), via the summary process 522, to the control process 520, and the control process 520 returns the summary $SUM||s_{SC}(D_{SUM})$ to the application process 500. The application process 500 receives the summary in step 654.

[0077] The individual signature and summary may be used by the application process 500, or any other process running on the host computer 100 in various ways outside the scope of this invention, including as proof of contract, for storage or for transmission to other entities.

[0078] Finally, in step 656, the control process 520 unfreezes the display, by recommencing receipt and processing of graphic primitives associated with the document image, and thereby in effect returns control of the display back to the application process 500 or other application software. Alternatively, control may not be handed back to the application process 500 until the user actuates the trusted switch 135 again, typically in response to another user message, which, this time, would not have a timeout period. This would give the user more time to review the static document image before returning the host computer to standard, non-trusted operation.

[0079] In order to verify a signed document, both the individual signature $PIX||s_{SC}(D_{PIX})||s_{SC}(FD)$ and the summary $SUM||s_{SC}(D_{SUM})$ must be verified. Such verification methods are well known to those skilled in the art of security. For example, the signature on the digest of the pixmap $s_{SC}(D_{PIX})$ is verified using the public key of the user, which is publicly available and preferably contained within a digital certificate $Cert_{SC}$ supplied by a certification authority. The verified digest is then compared with a value obtained by recomputing the digest from the pixmap, where the digest is generated using a standard, well-known and defined hash function. If the match is exact, the signature has been verified. Other signatures, including the summary, are checked in the same way.

[0080] A preferred method of enabling a person to verify the wording of the signed document is to translate the pixmap back into an image. This requires an application, or indeed a trusted display processor 260, to load the pixmap data PIX into the frame buffer memory 315 of a respective host computer 100. This allows a person to view the document that the signer has signed.

[0081] The stages of highlighting a document to be signed will now be described with reference to Figures 10a to 10d.

[0082] In the preferred embodiment, the seal data SEAL comprises a pixmap of a trusted image. For example, as shown in Figure 10a, the pixmap of the seal data 540 defines a 'smiley face' 1000. Figure 10b illustrates an image 1005 of an exemplary document Doc1 to be signed, in a window 1010 of the screen (not shown). As a first highlighting step, after the image has been made static but before the seal data has been received, the trusted display processor 260 highlights the document to be signed by superimposing a frame 1020 around the document image 1005, as illustrated in Figure 10c. Also, where a smartcard 122 is not present, a user message 1030 asking the user to insert his smartcard is displayed accompanied by a ten second countdown timer 1035, as also illustrated in Figure 10c. Next, when the smiley face pixmap image is retrieved from the smartcard 122, the trusted display processor 260 embellishes the frame 1040 with multiple instances 1045, or a mosaic, of the smiley face, as shown in Figure 10d. In addition, as shown in Figure 10d, the trusted display processor 260 generates a further user message 1050, accompanied by a ten second countdown timer 1055, asking the user whether they wish to proceed with the signing process. This embellished frame 1040 both indicates to the user that the correct static image area is being acted on and provides the user with a high level of confidence that the trusted display processor 260 is fully in control of the signing process; the presence of the user's own seal image provides confidence to the user that the message has come from the trusted display processor 260 rather than from some other (possibly subversive) software application or hardware device.

[0083] Figures 10e and 10g illustrate alternatives to the 'frame' visual effect illustrated in Figures 10c and 10d. In

Figure 10e, four single seal images 1060 are positioned at the corners of the static document image using the coordinates provided by the application process 500, to define the static image area. In Figure 10f, the static image is defined by modifying the background thereof to show a single seal image. In Figure 10g, the static image is defined by modifying the background thereof to show a mosaic of seal images. It is expected that the skilled reader will be able to think of other visual effects by which the static image may be highlighted in the light of the present description. In addition, it may be desirable to include further status messages during the signing operation, for example "Retrieving seal data 540 now...", "Generating document signature now...", etc.

[0084] It will be appreciated that the trusted display processor 260 needs to be able to display the seal image(s) and the messages in the correct places on screen. Clearly, the seal image and the message images are temporary, to the extent they appear during the signature process and disappear thereafter. There are well-known, standard display techniques for overlaying a first image with a second image, thereby obscuring a part of the first image, then removing the second image and restoring the portion of the first image that had been obscured. Such techniques are used as a matter of course in normal windows environments, for example, where multiple windows may overlap one another. The trusted display processor 260 is arranged to implement one of more of these standard techniques for the purposes of superimposing the seal image(s) and the message images over the standard display.

[0085] In some scenarios, it may be that a document is too large to fit all at once onto the VDU 105 screen and still be easily read by a person. Obviously, for the present embodiment to be practical, it is essential that a user can very clearly read the document before signing it. Therefore, the document can be split into multiple screen pages, each of which needs to be signed and cryptographically chained to the signature of the previous page, as will now be described.

[0086] First, the application process 500 causes the image of the first page to be displayed and makes a call to the trusted display processor 260 for signing as before. When the trusted display processor 260 returns the individual signature, instead of requesting a summary, the application process 500 instructs the trusted display processor 260 to display the image of the second page and sign the image. Clearly, in this case, the trusted display processor 260 is arranged to support such a request by the application process 500. Only after all images have been signed and returned to the application process 500 does the application process 500 issue a request for a summary. Then, the summary includes the number of images that were signed in this multi-page document, for example as illustrated in the two-page summary above.

[0087] The first page in the multi-page document is signed in the same way as a single page, resulting in return of an individual signature. When subsequent images are presented for signing, however, the trusted display processor 260 recognises that they are part of a multi-page document because no summary request was received after the previous signature request. As a result, the trusted display processor 260 displays a different message, which requests permission from the user to sign a continuation page. In response, the user who is signing a multi-page document uses the same reliable permission channel as before (for example, the trusted switch 135) to confirm to the trusted display processor 260 that this page is associated with the previous page, and is also to be signed. When the trusted display processor 260 receives this multi-page confirmation, it concatenates the signature of the previous signed page with the pixmap of the current page, creates a digest of the concatenation, and sends that to the smartcard for signing. This is instead of sending a digest of just the current pixmap. This process cryptographically 'chains' a subsequent page to the previous page, so that pages cannot be rearranged without detection, nor can intermediate pages be inserted or deleted without detection.

[0088] The validity of the first page may be checked in exactly the same way as a single page. The validity of subsequent pages is checked using the same method as for a single page, except that the digest of the current pixmap is replaced by the digest of the concatenated previous signature and current pixmap.

[0089] It will be appreciated that there are many ways of cryptographically chaining a subsequent page to a previous page. Such ways will be obvious to those skilled in the art of security in the light of the present description.

[0090] For added security, the image of each page of a multi-page document may be arranged to include the conventional footer 'Page x of y, where 'x' is the number of the page and 'y' is the total number of pages. This enables ready detection by a person of a truncated document simply by reading the document.

[0091] A significant benefit of the present document signing scheme is that a signed document can be re-signed and countersigned. As such, it is preferable for the summary of a document to include an audit trail. There are many variations on re-signing and countersigning, although (obviously) an electronic integrity check should always be done before any further signing. At one extreme, the new signer could view, confirm and re-sign each signed image in turn, effectively replacing the original signature by a new one. This method could be used, for example, by a user signing a document prepared for him by someone else. At the other extreme, the new signer could simply 'rubber stamp' the original signature by signing the original summary, without necessarily viewing the document at all. This could be useful to a manager countersigning the work of a trusted employee.

[0092] For a re-signing operation, the application process 500 issues a re-signing request, and transmits an already signed document (plus the individual signature(s) and the summary) to the trusted display processor 260. The trusted display processor 260 verifies the signed document using the public key of the signer, recovers the pixmap of the

document (or each page of the document) and displays each verified image in the correct order to the new user, as if they were original images from the signature request application. The user confirms the acceptance of each individual image, for example using a trusted switch 135 as before, and causes the images to be signed as before by a smartcard belonging to the new user. This results in a signed document that is the same as the original document, except that it has been signed by the smartcard belonging to the new user.

[0093] Similarly, for a counter-signing operation, the application process 500 issues a counter-signing request and transmits the signed document (plus individual signatures and the summary) to the trusted display processor 260. The trusted display processor 260 verifies the signed document and displays each verified image in the correct order to the new user, as if they were original images from the application process 500. The user confirms acceptance of each individual image and the trusted display processor 260 signs the original summary using the smartcard belonging to the new user. Optionally the new user could provide a certificate of the previous user's public key, signed by the new user, to ease the processing overhead associated with later verification of the signature.

[0094] Clearly, there are many possible variations on the theme of re-signing and countersigning, which will be apparent to the skilled person in the light of the present description.

[0095] Since a document may have a history of signing, re-signing and/or counter-signing, the present embodiment conveniently provides audit information, which forms part of the document summary. This audit information allows the signature history of the document to be traced. The audit information includes data about the previous state of the document and the actions taken by the new user to create the new state of the document. The audit information is signed by the trusted display processor 260, since the audit information must be independent of the user. The audit information always contains any previous summary information (including the signature on that summary information, by the previous signer). If the signed document has been created from scratch, the identity label l_{DP} of the trusted display processor 260 is inserted as an audit root. The audit information preferably also includes an indication of which individual images were viewed and confirmed by the new user, and whether the document was created from scratch, or was re-signed, or was countersigned by the new user. To create a summary including audit information, the smartcard is sent a digest of the audit information concatenated with the previously described contents of a summary, rather than a digest of just the previously described contents of a summary. The rest of the process is as previously described.

[0096] An enhancement to the process for signing a document is that, prior to signing the pixmap data, the trusted display processor 260 compresses the pixmap using a lossless compression algorithm so that the overheads associated with storing and sending the individual signature are reduced.

[0097] The pixmap may be compressed by standard compression algorithms, for example a codeword-based algorithm applying LZ-1 or LZ-2 compression. Alternatively, a technique similar to OCR (optical character recognition) may be used to compress the pixmap. In this case, the situation differs from conventional OCR in that the input data has been perfectly 'scanned', albeit at a lower resolution than in conventional OCR. The OCR-compressed version of the pixmap may be generated by 'blob-matching' to create an alphabet for the pixmap, constructing a pixmap of each character in the alphabet, and constructing a message using those characters, such that the message represents the original pixmap. This means that the pixmap can be compressed to a new alphabet and a message written in that alphabet. Since there are, obviously, no errors nor ambiguity in the pixmap data, this is a lossless compression method.

[0098] Another way of reducing the size of the image pixmap is by representing the image as a pure black and white image, requiring only a single bit - set to zero or one - to define whether a pixel is black or white. Otherwise, the document image is represented as a full colour image, where each pixel may typically require up to 24-bits. Obviously, this technique may be suitable for simple, black and white text-based documents. However, it would not be appropriate for colour documents or images.

[0099] At any time, the document image may be converted back into a text-based document using an OCR-type process to reconstruct a standard digital textual representation of the document. This technique cannot be used in the signature, since the textual mapping may be incorrect, but can be used by the receiver of a signed document to convert it back into a standard digital textual representation (such as ASCII) for subsequent machine manipulation. In preferred embodiments, the trusted display processor 260 is equipped to enact OCR document recovery.

[0100] To enact OCR, an OCR alphabet is generated in a standard fashion and is then matched to stored fonts and hence converted to a standard character set. As in conventional OCR, ambiguous matches may be retained as a pixmap and flagged for conversion by the user. (This is unlikely, particularly if font type and size information has been supplied in the display format data FD, because there is no error in the data.) In cases of extreme caution, the entire reconstructed document should be manually checked by a person against the view of the document that the signer intended to sign.

[0101] Preferably all document reconstruction processes are done by processes that are trusted.

[0102] The preferred embodiment described above relies on the premise that the trusted display processor 260 has direct and exclusive access to video data stored in the frame buffer memory 315, beyond the point where the video data can be manipulated by host computer 100 software, including the operating system. This implies that the video data cannot be modified unless the trusted display processor 260 makes the modification.

[0103] It will be appreciated that not all computer architectures are arranged in this way. For example, some computer architectures are arranged such that the frame buffer memory forms a part of the main memory, thus forming a single address space (SAS) display system. One benefit of such a system is that both the CPU and the display processor can access the frame buffer memory and share the graphics operation overhead, thereby improving graphics performance. Clearly, an implementation of the present invention in such a SAS system cannot rely on the premise that the buffer memory is safe during signing, since the CPU can still access the memory. However, there are many ways in which such a SAS system may be modified to support implementations of the present invention. For example, the memory could be provided with a control line from a trusted display processor such that, during a signing operation, the memory is prevented from being updated by data from the CPU. The memory devices themselves are preferably modified so that they include the extra logic to perform this function. Alternatively, access to memory is blocked by other logic circuits inserted into the normal control path of the memory. Such systems, therefore, rely on the modified premise that the video data in the frame buffer memory can only be modified, other than by the trusted display processor, with the permission of the trusted display processor. Clearly, this premise is as valid for secure operation as the first premise, as long as the system is truly secure.

[0104] In other architectures, for example in simple graphics environments, the functionality of a display processor may form part of the operating system itself, thereby removing the requirement for separate display processor hardware. Clearly, in this case, the graphics overhead put on the CPU will be higher than in a system with separate display processor hardware, thereby limiting the graphics performance of the platform. Clearly, there is then no place for a 'trusted display processor' as such. However, it will be apparent to the skilled person that the same function as provided by the trusted display processor, that of protecting the frame buffer memory and interacting with a smartcard, can be implemented using an appropriate trusted component, which controls the display system (in whatever form) during signing.

[0105] In other embodiments of the invention, in addition or alternatively, the trusted display processor (or equivalent) includes an interface for driving a trusted display. The trusted display might be, for example, an LCD panel display. In the same way that the trusted switch provides a trusted means for a user to interact with the trusted display processor, the trusted display can provide a trusted means for feeding back information to the user other than via the standard VDU. For example, the trusted display might be used to provide user status messages, as described above, relating to a signing operation. As such, applications running on the standard host computer should not be able to access the trusted display, because the display is connected either directly to the trusted display processor or via some form of trusted channel. In essence, such a trusted display is an addition to the so-called 'trusted interface' described above. In practice, there is no reason why other forms of trusted feedback device, of which the trusted display is one example, could not be included in addition, or as an alternative. For example, there may be scenarios where some form of trusted sound device would be useful for providing audible feedback.

Claims

1. A data processing system arranged to generate a digital signature representative of a document, the data processing system comprising;

main memory means for storing a document to be digitally signed;
main processing means for executing at least one application process comprising means to generate graphics signals for displaying the document;
means for generating a request signal for the document to be signed;
a display system comprising:

frame buffer memory;
means for generating digital image data representative of the document on the basis of the graphics signals and storing the digital image data in the frame buffer memory; and
means for reading the digital image data from the frame buffer memory, converting the data into signals suitable for displaying an actual image thereof on a display means and forwarding said signals to a display means; and

a trusted component comprising independent processing means operable, in response to receipt of the request signal, for generating a digital signature representative of the digital image data.

2. A data processing system according to claim 1, wherein the trusted component comprises means for denying to any unauthorised application or process write access to at least the portion of the frame buffer memory containing

the digital image data of the document, and means for generating a digital signature representative of the digital image data while the respective portion of the frame buffer memory is not accessible for writing data to.

- 5 3. A data processing system according to claim 1 or claim 2, further comprising a removable token comprising processing means for receiving the digital image data, or a representation thereof, and generating a respective digital signature.
- 10 4. A data processing system according to any one of the preceding claims, wherein the trusted component further comprises means for acquiring and/or generating trusted image data and means for controlling the display system to highlight the displayed document image using the trusted image data.
- 15 5. A data processing system according to claim 4, wherein the trusted image data comprises pixmap data representative of the trusted image or instructions for forming the trusted image.
- 20 6. A data processing system according to claim 4 or claim 5, wherein the trusted component controls the display system to highlight the displayed document image by producing one or more of the following visual effects:
 - a border, or an indicator or indicators defining a border, characterised by the trusted image and positioned at least partly around the document image;
 - 20 a background pattern characterised by the trusted image forming at least part of the background of the document image;
 - an image characterised by the trusted image formed within the document image; and/or
 - a text message characterised by the trusted image formed within or near the document image.
- 25 7. A data processing system according to any one of claims 4 to 6, wherein the trusted component comprises means for acquiring and/or generating trusted image data from a removable token.
- 30 8. A data processing system according to any one of the preceding claims, wherein the trusted component further comprises means for controlling the display system to display messages to a user.
- 35 9. A data processing system according to claim 8, further comprising trusted input means by which a user can respond to messages in a secure fashion.
- 40 10. A data processing system according to claim 9, wherein the trusted input means comprises a switch connected to the trusted component via a secure communications channel.
- 45 11. A data processing system according to claim 3 or claim 7, wherein the trusted component and the secure token enact a mutual authentication process in advance of further interactions.
- 50 12. A data processing system according to any one of the preceding claims, wherein the trusted component forms an integral part of the display system.
- 55 13. A data processing system according to claim 12, wherein the display system is arranged such that the trusted component is physically and functionally positioned between the main processing means and the frame buffer memory, such that the main processing means can only access the frame buffer memory indirectly through functions of the trusted component.
14. A data processing system according to any one of the preceding claims, further comprising means for generating data summarising a digital signature operation.
15. A method for digitally signing a document, comprising the steps:
 - generating digital image data of the document and updating the digital image data in a frame buffer memory;
 - reading the digital image data from the frame buffer memory, converting the digital image data into signals suitable for driving a visual display means and transmitting the signals to a visual display means for displaying an image of the document; and
 - on demand, reading the digital image data from the frame buffer memory and generating a digital signature representative thereof.

16. A method according to claim 15, further comprising the step of temporarily denying write access to the frame buffer memory by unauthorised processes while generating the digital signature.
- 5 17. A method according to claim 14 or claim 15, further comprising the step of acquiring and/or generating trusted image data and using the trusted image data to highlight the document image.
18. A method according to claim 17, wherein step of highlighting the document image is achieved by generating any one or more of the following visual effects:
 - 10 a border, or an indicator or indicators defining a border, characterised by the trusted image and positioned at least partly around the document image;
 - a background pattern characterised by the trusted image forming at least part of the background of the document image;
 - an image characterised by the trusted image formed within the document image; and/or
 - 15 a text message characterised by the trusted image formed within or near the document image.
19. A method according to claim 17 or claim 18, wherein the trusted image data is acquired from a removable token.
20. A data processing system arranged to digitally sign a document in accordance with any one of claims 15 to 19.
- 20 21. A method for digitally signing a document comprising a plurality of individual viewable pages, comprising the steps:
 - a) generating digital image data of the first page of the document and updating the digital image data in a frame buffer memory;
 - 25 b) reading the digital image data from the frame buffer memory, converting the digital image data into signals suitable for driving a visual display means and transmitting the signals to a visual display means for displaying an image of the document;
 - c) reading the digital image data from the frame buffer memory, generating a digital signature representative thereof and storing the digital signature;
 - 30 iv) repeating steps a) to c) for the other page(s) of the document; and
 - v) generating a further digital signature representative of all previous digital signatures.
22. A method for a second user to digitally counter-sign a document that has already been signed by a first user, the document being accompanied by a respective first digital signature generated by using a secret of the first user,
35 comprising the steps:
 - generating digital image data of the document and updating the digital image data in a frame buffer memory;
 - reading the digital image data from the frame buffer memory, converting the digital image data into signals suitable for driving a visual display means and transmitting the signals to a visual display means for displaying
40 an image of the document;
 - verifying the integrity of the first digital signature; and
 - on the basis of a secret of the second user, generating a digital signature representative of the first digital signature.
- 45 23. A method for a second user to digitally re-signing a document that has already been signed by a first user, the document being accompanied by a respective first digital signature generated by using a secret of the first user, comprising the steps:
 - generating digital image data of the document and updating the digital image data in a frame buffer memory;
 - 50 reading the digital image data from the frame buffer memory, converting the digital image data into signals suitable for driving a visual display means and transmitting the signals to a visual display means for displaying an image of the document;
 - verifying the integrity of the first digital signature; and
 - reading the digital image data from the frame buffer memory and, on the basis of a secret of the second user,
55 generating a digital signature representative of the first digital signature.
24. A data processing system arranged to generate a digital signature representative of a document, the data processing system comprising;

main processing means for generating graphics signals for displaying a document;
a display system comprising:

frame buffer memory;

means for generating digital image data representative of the document on the basis of the graphics signals and storing the digital image data in the frame buffer memory; and

means for reading the digital image data from the frame buffer memory, converting the data into signals suitable for displaying an actual image thereof on a display means and forwarding said signals to a display means; and

means for generating a digital signature representative of the digital image data.

25. A system for digitally signing a document comprising:

frame buffer memory;

means for generating image data representative of a document and storing the image data in the frame buffer memory;

means for reading the image data from the frame buffer memory and displaying a respective image on a display means; and

means for reading the image data from the frame buffer memory and generating a digital signature thereof.

26. A trusted component for use in a data processing system according to any one of claims 1 to 14.

27. A trusted component according to claim 26, fabricated to be tamper resistant.




FIGURE 1






FIGURE 2

**FIGURE 3****FIGURE 4**

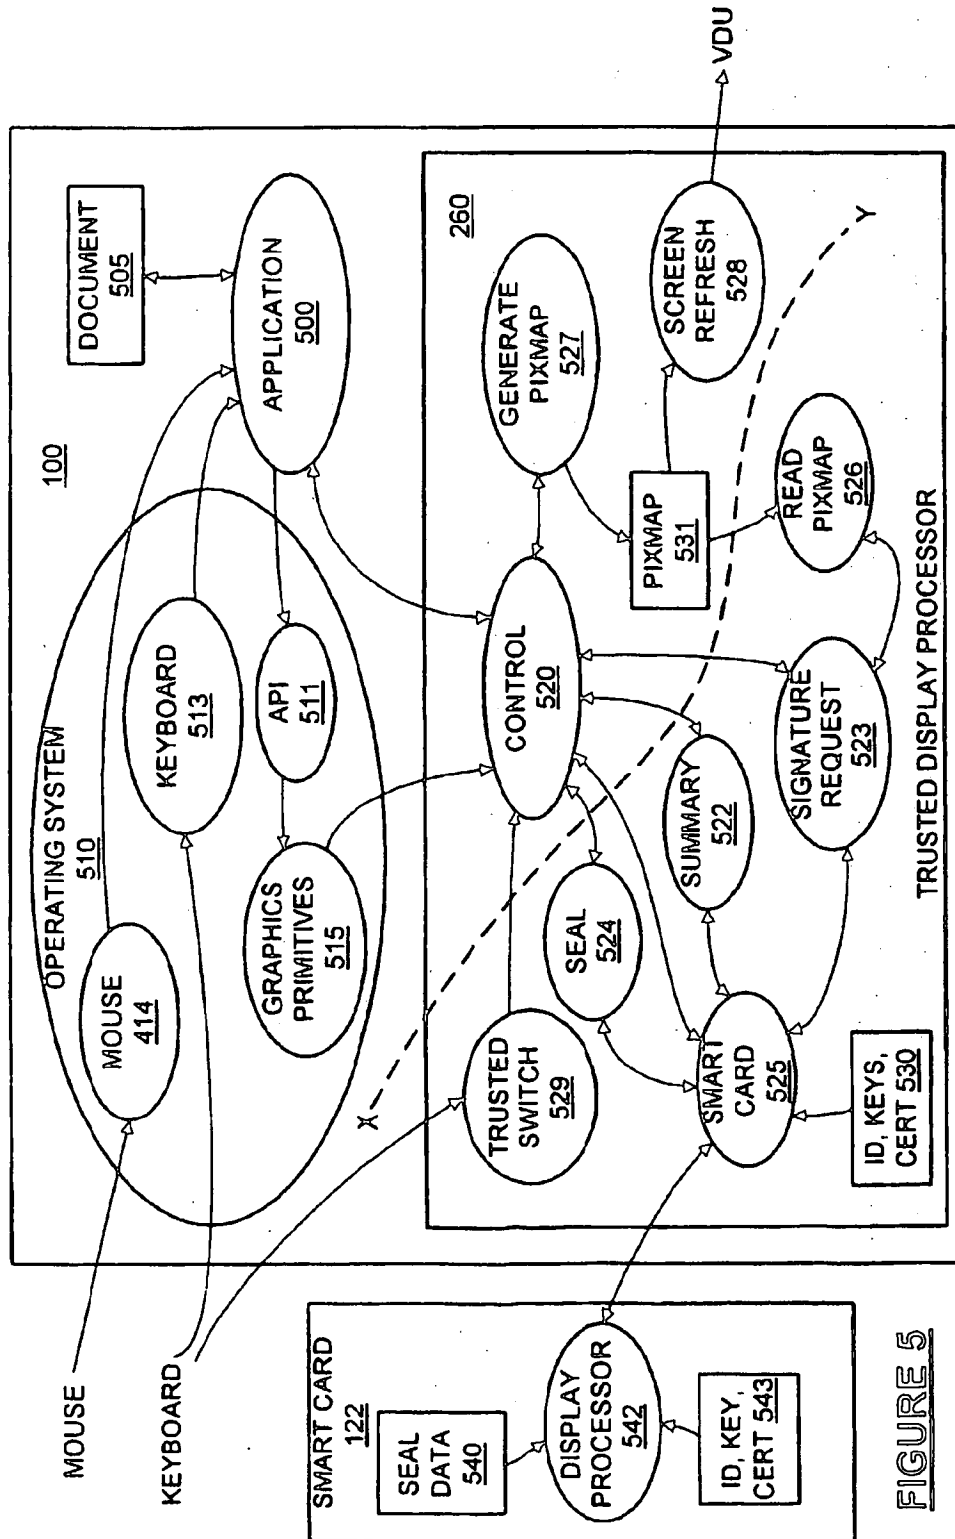
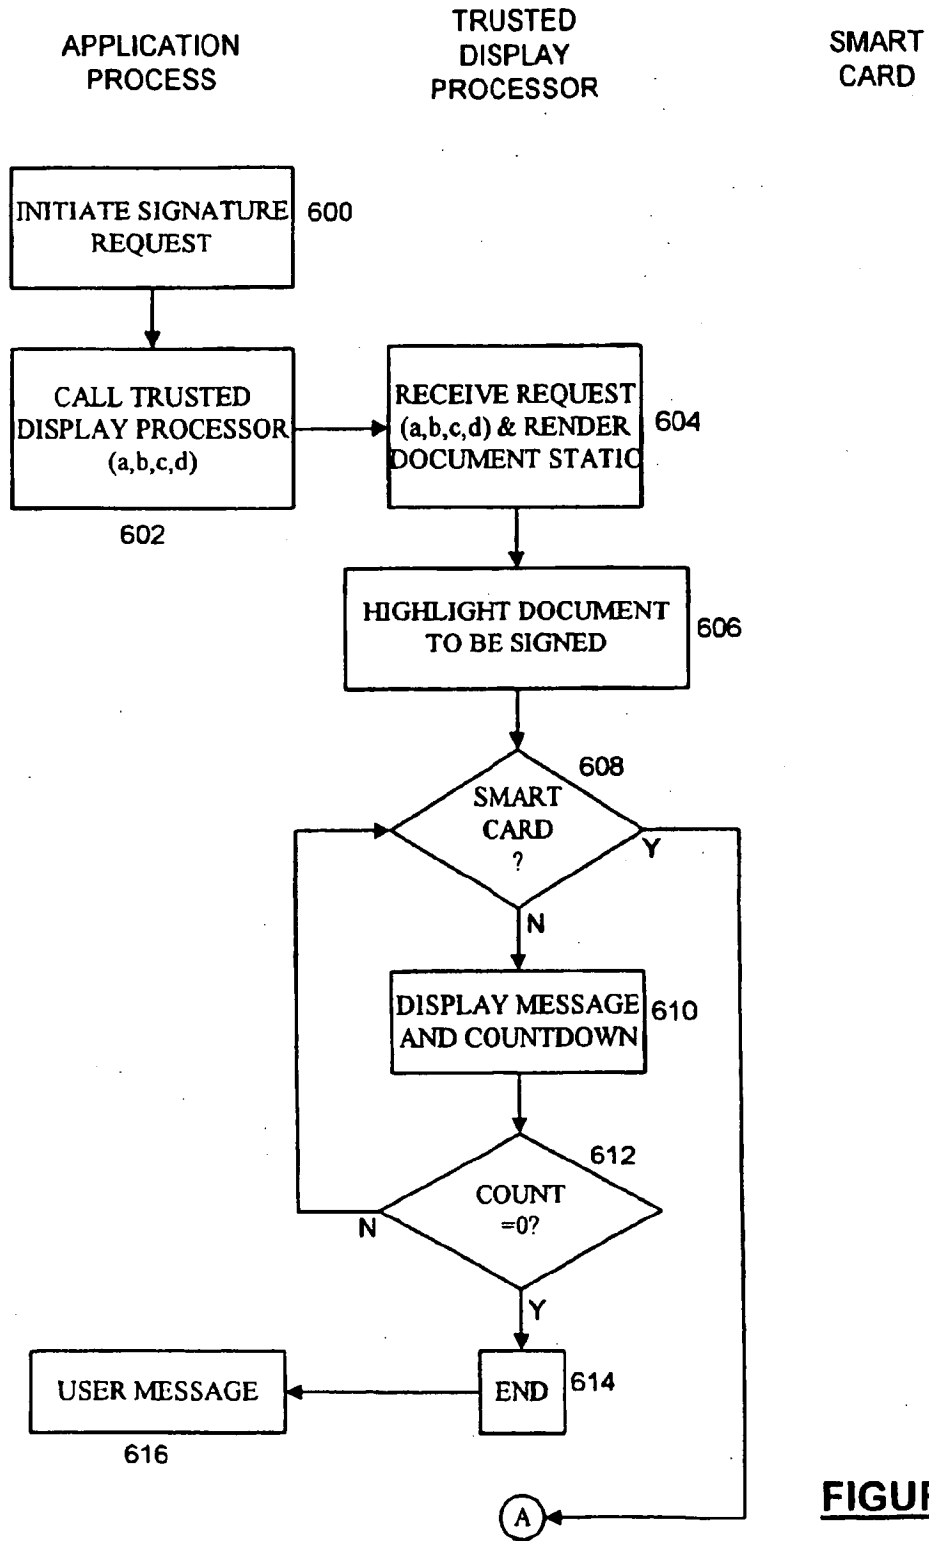
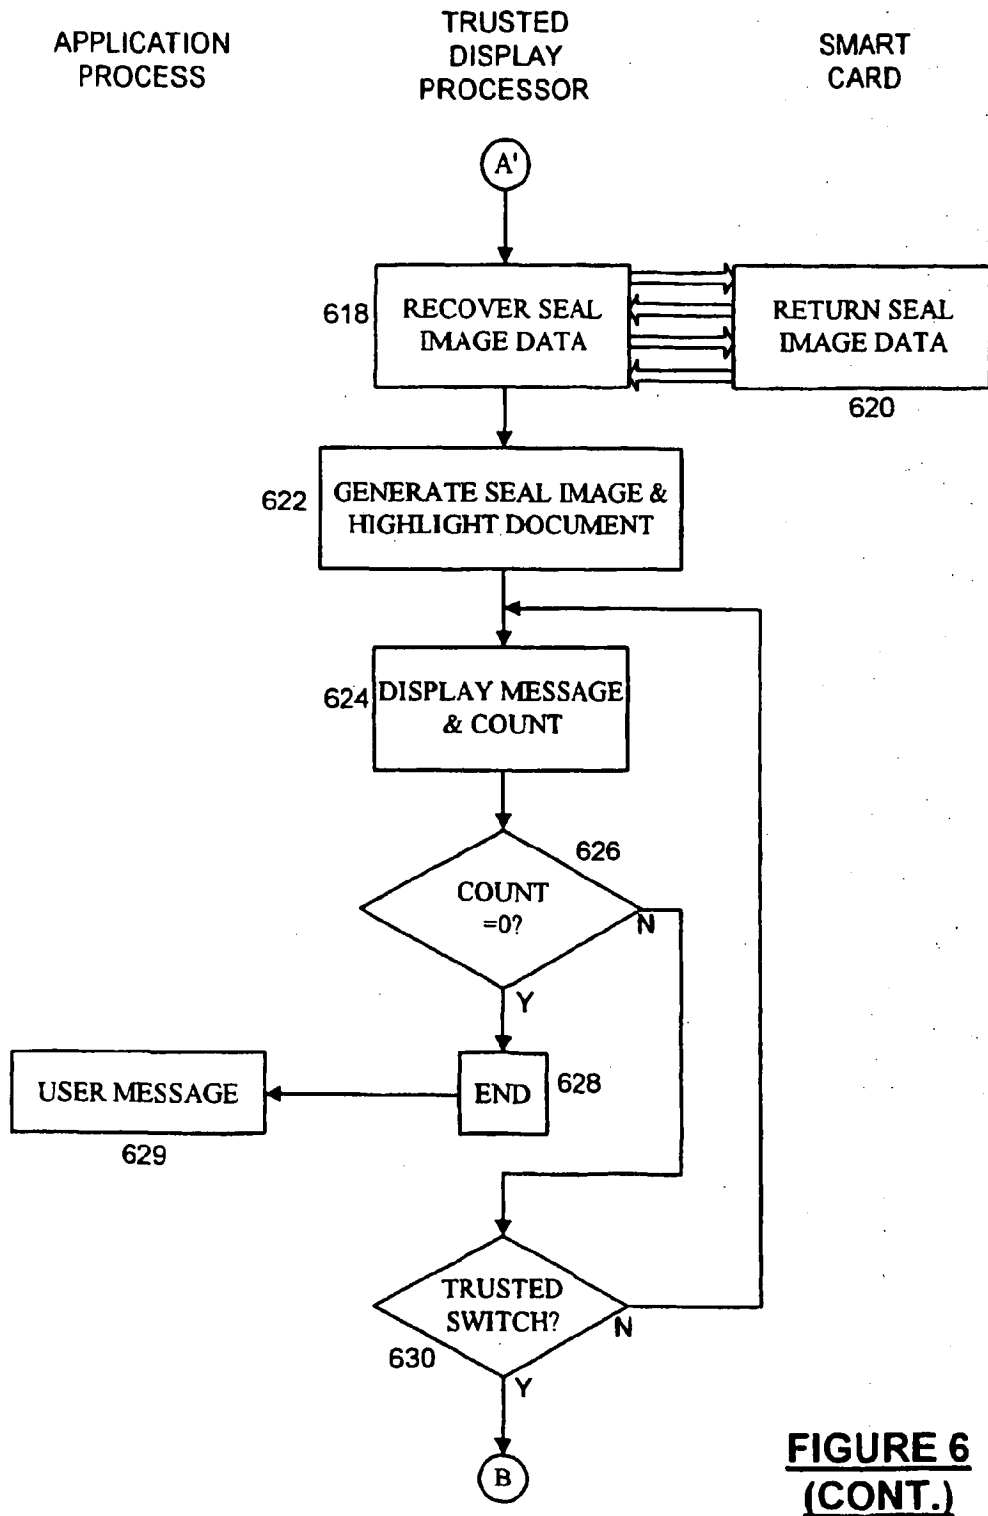


FIGURE 5

**FIGURE 6**

**FIGURE 6
(CONT.)**

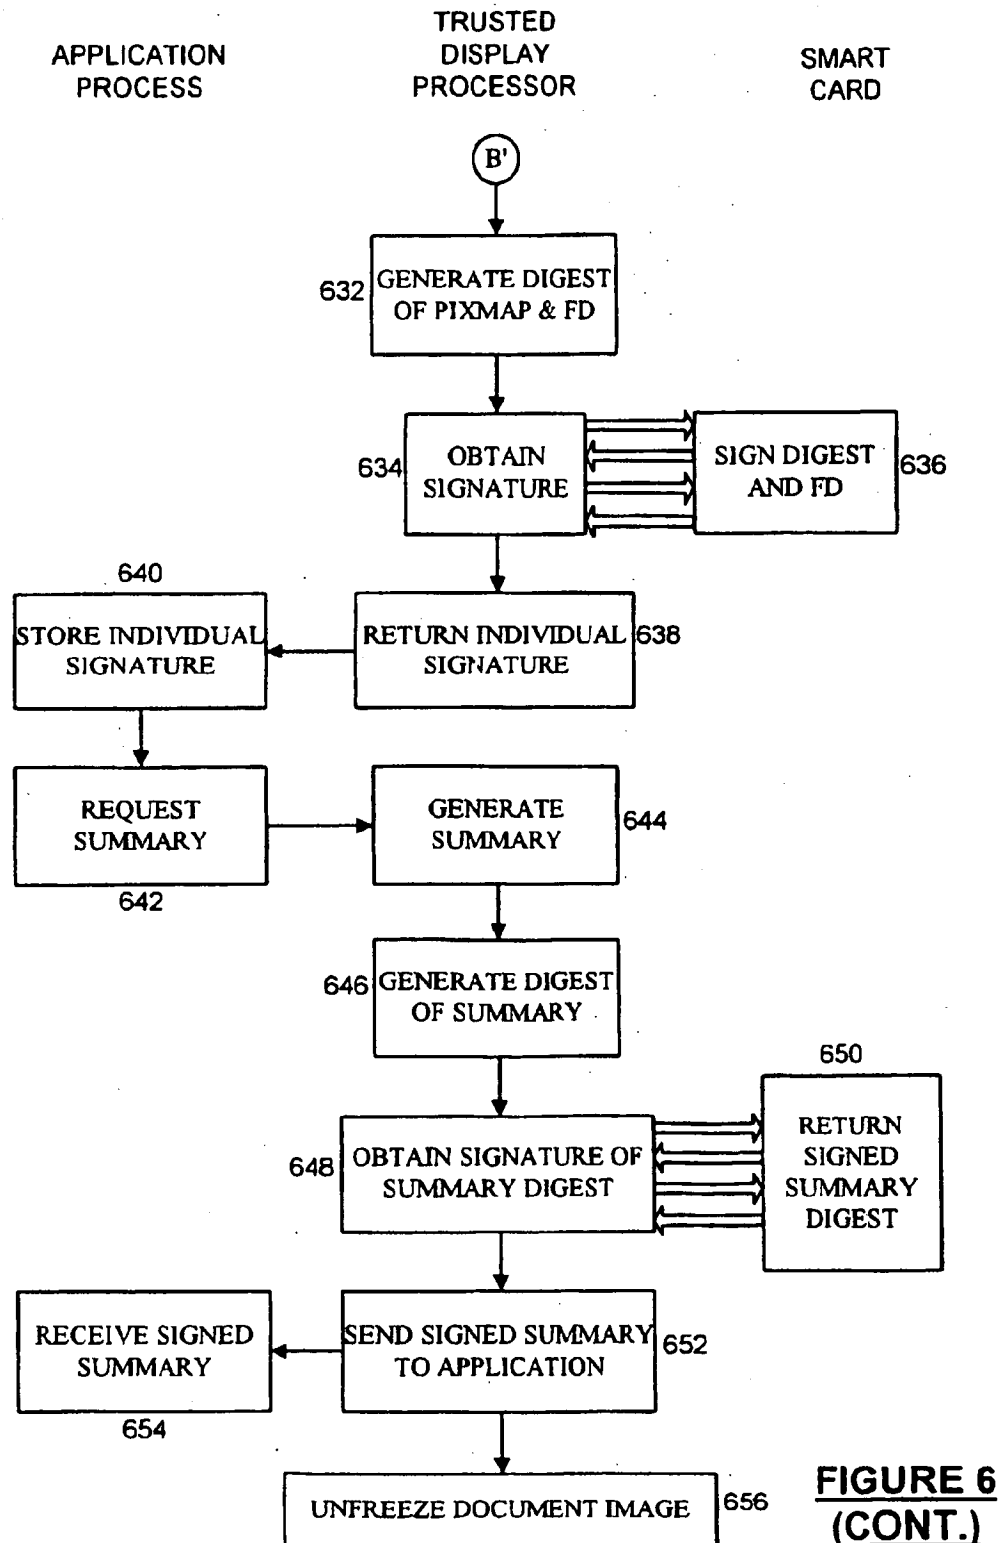


FIGURE 6
(CONT.)

TRUSTED DISPLAY
PROCESSOR
(SMART CARD PROCESS)

SMART CARD
(DISPLAY PROCESSOR
PROCESS)




FIGURE 7




FIGURE 8




FIGURE 9




FIGURE 10a




FIGURE 10b




FIGURE 10c




FIGURE 10d




FIGURE 10e




FIGURE 10f




FIGURE 10g

EUROPEAN SEARCH REPORT

Application Number
EP 99 30 4164

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
X	EP 0 717 337 A (IBM) 19 June 1996 (1996-06-19) * column 1, line 47 - column 2, line 40 * * column 14, line 25 - line 40 *	15,20, 24,25	G06F1/00
Y		1,3, 8-13, 16-18, 21-23, 26,27	
X	US 5 907 619 A (DAVIS DEREK L) 25 May 1999 (1999-05-25) * the whole document *	15,20, 24,25	
A		1,21-23, 26,27	
Y	WOO T Y C ET AL: "AUTHENTICATION FOR DISTRIBUTED SYSTEMS" COMPUTER,US,IEEE COMPUTER SOCIETY, LONG BEACH., CA, US, vol. 25, no. 1, 1 January 1992 (1992-01-01), pages 39-52, XP000287833 ISSN: 0018-9162 * figure 2 * * page 45, left-hand column, paragraph 2 - page 48, left-hand column, last paragraph *	1,3,11, 26,27	
			TECHNICAL FIELDS SEARCHED (Int.Cl.7)
			G06F
Y	WO 94 01821 A (SECURE COMPUTING CORP) 20 January 1994 (1994-01-20) * abstract; figures 2,3,5,6 * * page 9, line 8 - page 11, line 18 * * page 13, line 33 - page 14, last line * * page 16, line 6 - page 17, line 28 *	8-10,12, 13	
A		1,2,4, 17,18	
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 29 February 2000	Examiner Powell, D
CATEGORY OF CITED DOCUMENTS		T: theory or principle underlying the invention E: earlier patent document, but published on, or after the filing date D: document cited in the application L: document cited for other reasons X: particularly relevant if taken alone Y: particularly relevant if combined with another document of the same category A: technological background O: non-written disclosure P: intermediate document	

European Patent
Office

EUROPEAN SEARCH REPORT

Application Number

EP 99 30 4164

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.CL7)
Y	EP 0 386 867 A (FISCHER ADDISON M) 12 September 1990 (1990-09-12) * the whole document *	21-23	
Y	WISEMAN S ET AL: "THE TRUSTED PATH BETWEEN SMITE AND THE USER" PROCEEDINGS OF THE SYMPOSIUM ON SECURITY AND PRIVACY, US, WASHINGTON, IEEE COMP. SOC. PRESS, vol. -, 1988, pages 147-155, XP000011725 * page 152, left-hand column, line 1 - last line *	16-18	
A		2,4,6,8,9	
A	BERGER J L ET AL: "COMPARTMENTED MODE WORKSTATION: PROTOTYPE HIGHLIGHTS" IEEE TRANSACTIONS ON SOFTWARE ENGINEERING, US, IEEE INC. NEW YORK, vol. 16, no. 6, 1 June 1990 (1990-06-01), pages 608-618, XP000128949 ISSN: 0098-5589 * page 615, left-hand column, paragraph 4 - page 616, right-hand column, paragraph 1 *	2,4,6,17,18	
The present search report has been drawn up for all claims			TECHNICAL FIELDS SEARCHED (Int.CL7)
Place of search THE HAGUE		Date of completion of the search 29 February 2000	Examiner Powell, D
CATEGORY OF CITED DOCUMENTS X: particularly relevant if taken alone Y: particularly relevant if combined with another document of the same category A: technological background O: non-written disclosure P: intermediate document T: theory or principle underlying the invention E: earlier patent document, but published on, or after the filing date D: document cited in the application L: document cited for other reasons &: member of the same patent family, corresponding document			

EPO FORM 1503 06/92 (PUB01)

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 99 30 4164

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

29-02-2000

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 0717337 A	19-06-1996	JP 8221268 A	30-08-1996
		US 5745678 A	28-04-1998
US 5907619 A	25-05-1999	NONE	
WO 9401821 A	20-01-1994	US 5596718 A	21-01-1997
		AU 663406 B	05-10-1995
		AU 4672693 A	31-01-1994
		EP 0649546 A	26-04-1995
		JP 7509086 T	05-10-1995
		US 5822435 A	13-10-1998
EP 0386867 A	12-09-1990	US 5005200 A	02-04-1991
		AT 113429 T	15-11-1994
		AT 150605 T	15-04-1997
		AU 620291 B	13-02-1992
		AU 4242589 A	13-09-1990
		CA 2000400 A, C	07-09-1990
		DE 69013541 D	01-12-1994
		DE 69013541 T	09-03-1995
		DE 69030268 D	24-04-1997
		DE 69030268 T	26-06-1997
		DK 386867 T	03-04-1995
		EP 0586022 A	09-03-1994
		ES 2036978 T	01-01-1995
		ES 2098651 T	01-05-1997
		GR 93300050 T	30-06-1993
		JP 2291043 A	30-11-1990
		US 5214702 A	25-05-1993

EPO FORM P/046

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82